

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

Patent Application for

Scalable Policy Server

Invention of: Kwok Ho Chan
26 Woods Court
Dunstable, MA 01827-2103

Attorney docket number:
2204/A87
12541BAUS01

Attorneys:
Bromberg & Sunstein LLP
125 Summer Street
Boston, MA 02110-1618
Tel: (617) 443-9292
Fax: (617) 443-0004

Scalable Policy Server

This application claims priority from U.S. provisional
5 patent application 60/205,706, filed May 19, 2000, which is
incorporated herein by reference.

Field of the Invention

The invention generally relates to computer networks
10 and, more particularly, to establishing communications
across a network.

Background Art

Network devices communicate across a network (e.g., the
15 Internet) in accord with many known protocols. One
protocol, known as the Resource Reservation Protocol (the
"RSVP protocol"), enables one or more network devices to
reserve bandwidth based upon some pre-specified network
policy. The policy may be based upon any desired
20 characteristic, such as the identity of some given user
(e.g., an executive of a corporation requesting data) or
network device (the prior noted executive's computer), or
the type of data being transmitted (e.g., video on demand).

When utilizing the RSVP protocol, a sending access
25 device (e.g., a personal computer) first transmits a "path
message" toward a receiving access device across the
network. Among other things, the path message typically
includes a request for a specified amount of bandwidth
(e.g., a Quality of Service level), and some policy data
30 (e.g., identity of sender, type of data, etc...). A network
device, such as a router, between the two access devices
then receives the path message, and communicates with a
policy server to determine some preliminary policy issues

for the two access devices. For example, the policy server may indicate the broadest policy parameters that the two access devices may utilize. The router then forwards the path message to the receiving access device.

5 Upon receipt of the path message, the receiving access device responsively generates and transmits a reservation message toward the sending access device. Among other data, the reservation message typically includes data indicating the amount of bandwidth that the receiving access device can
10 handle (within the amount requested in the path message), and other policy data. The router receives the reservation message before the sending access device receives such message. The router then communicates with the policy server once again, requesting that the policy server apply
15 the policy to the information in the two messages. The policy server responsively applies the policy, and then forwards such applied policy to the router. Upon receipt from the policy server, the router forwards the applied policy to the sending access device. The sending access
20 device then utilizes the applied policy to mark specified data packets transmitted between the two access devices. Intermediate network devices thus may process such data packets in accord with the applied policy.

 In commonly used systems, one policy server is utilized
25 with many routers. Accordingly, problems can arise when many users are indirectly utilizing a single policy server. Specifically, processing both the path message and the reservation message, as well as applying the policy to the specified data, can degrade network performance.

Summary of the Invention

In accordance with one aspect of the invention, a method and apparatus for establishing communication (across a network) between a first network device and a second network device apply policy data received from a policy device (e.g., a policy server), thus offloading significant functionality from the policy device. In illustrative embodiments, this policy application is executed by a third network device. Offloading such functionality reduces the processing required by the policy device, thus improving network performance. To that end, the method and apparatus store policy data in a storage device that is accessible by the third network device. The policy data is provided by the policy device in response to a communication request message from the first network device to the second network device, and is specifically related to communication between the first network device and the second network device. The third network device receives a confirmation message sent by the second network device in response to the communication request message. The confirmation message indicates that the second network device is prepared to receive communications from the first network device. The stored policy data and the confirmation message are forwarded from the third network device to the first network device.

In representative embodiments, the communication request may be a path message, and the confirmation message may be a reservation message. Such messages thus are transmitted in accord with the well known Resource Reservation Protocol (the "RSVP protocol"). In such case, the policy device may be a policy server, the first network device may be a network access device (e.g., a computer, network appliance, cell phone, or other device that permits

a user to access a network), and the network may be the Internet. In addition, the third network device may be, for example, a router.

A representative embodiment further marks data traffic
5 directed from the first network device to the second network device according to the stored policy. This data traffic may include, without limitation, User Datagram Protocol (UDP) or Transmission Control Protocol (TCP) traffic, or both. The data traffic may also include streaming video or
10 audio. The third network device may be remote from the policy device, and may include a computer system with local storage that includes the storage device.

In a specific embodiment, the stored policy data determines a Quality of Service (QoS) characteristic (e.g.,
15 reserved bandwidth) that is assigned to the communications from the first network device to the second network device. The policy device may use the Common Open Policy Service (COPS) protocol. The stored policy data may also include either a Differentiated Services Code Point (DSCP) or 802.1p
20 signaling marking, or both.

Brief Description of the Drawings

The present invention will be more readily understood by reference to the following detailed description taken
25 with the accompanying drawings, in which:

Figure 1 schematically shows a representative network arrangement that may be utilized to implement various embodiments of the invention.

Figures 2A and 2B show a typical process of
30 establishing communication between the two access devices shown in Figure 1.

Detailed Description of Specific Embodiments

Figure 1 schematically shows an exemplary network arrangement that may be utilized with illustrative embodiments of the invention. In particular, the network arrangement includes a sending access device 10 that is attempting to establish a communication link/channel (hereinafter "link") with a receiving access device 12 across a large scale, public or private network 14 (e.g., the Internet or an intranet). For example, the sending access device 10 may be attempting to establish a link for transmitting streaming video data to the receiving access device 12. The access devices 10 and 12 may be any type of device that is capable of accessing a network 14, such as a computer system, server, Internet appliance, modem, cellular telephone, wireless access product, local area network access product, virtual private network access product, or Internet service provider access product. Note that this listing is exemplary and not intended to limit the various embodiments of the invention. It also should be noted that the access devices 10 and 12 may be referred to as "network devices."

Among other things, the network 14 includes a policy server 16 for facilitating policy driven processes (discussed below), and first and second routers 18 and 20. Illustrative embodiments implement the well known Resource Reservation Protocol (the "RSVP protocol") over the Internet Protocol (the "IP protocol"). It should be noted, however, that although discussed in terms of the RSVP protocol and the IP protocol, illustrative embodiments may be applicable to other similar protocols. Accordingly, discussion of IP and RSVP are not intended to limit the various embodiments of the invention.

Figures 2A and 2B show a process of establishing a communication link between the sending and receiving access devices 10 and 12. Such link is a one-way link from the sending access device 10 to the receiving access device 12 for transmitting one designated information flow. Two-way data transmission can be executed between the two access devices 10 and 12, however, if each access device 10 and 12 executes the process for its respective transmission flows. An information flow is defined by the application generating the data to be transmitted. For example, in a video conference between the two access devices 10 and 12, the video data can be considered a first flow, while the voice data can be considered a second flow. Accordingly, each access device 10 and 12 executes the process one time for the first (video) flow, and a second time for the second (voice) flow. As a further example, both the video data and voice data of the previous example can be considered a single information flow.

In illustrative embodiments of the process, the two access devices 10 and 12 attempt to reserve a specified bandwidth (complying with well known Quality of Service ("QoS") processes) for transmitting an information flow from the sending access device 10 to the receiving access device 12. Accordingly, if successfully established, such specified bandwidth is "guaranteed" to the two access devices 10 and 12 when they communicate. The process begins at step 200, in which an RSVP path message is generated and transmitted by the sending access device 10 toward the receiving access device 12 via the first router 18. In illustrative embodiments, the path message signals a QoS requirement to the network 14. To that end, the path message may include:

- a session object describing the receiving access device 12 via a receiver access device IP address, port number, and protocol type;
- a sender template object describing the sending
5 access device 10 via the sender access device IP address and port number; and
- a policy object that carries policy information. This object can carry many different types of information. For example, such information can include data regarding
10 inter-operability with the well-known Microsoft Windows 2000™ operating system (e.g., user ID and application ID information).

Once the first router 18 receives the path message, the process continues to step 202, in which the path message is
15 forwarded to the policy server 16. For example, if the path message is new (i.e., it is the first path message for the desired information flow) or has changed (i.e., it is different from a previous path message received for the desired information flow), then the first router 18 notifies
20 the policy server 16 as such by utilizing a Common Open Policy Service Protocol request message (a "COPS" request message). Specifically, a new COPS handle is used for a new path message, while an existing COPS handle is used for a changed path message. The COPS request message carries the
25 path message to fully communicate the specific data to the policy server 16.

Once it receives the COPS request message, the policy server 16 determines and applies the relevant policy, and forwards such applied policy to the first router 18 (step
30 204). Specifically, the policy server 16 determines:

- the receiving access device 12 from the receiver IP address, receiver port number, and protocol type in the session object;

- the sender access device 10 from the sender IP address and sender port number in the sender template object;

- the specific user of the sending access device 10 from the user ID information in the policy object; and

- specific application and sub-application data from the application ID information in the policy object.

Based upon this and other information, the policy server 16 makes its policy decision and selects a DSCP (differentiated services code point) and/or 802.1p priority marking. Accordingly, as noted above with regard to step 204, the policy server 16 forwards a COPS decision message with a DSCP and/or 802.1p marking to the first router 18.

Upon receipt of the COPS decision message, the first router 18 installs the DSCP and/or 802.1p marking information in its path state (step 206). Stated another way, the applied policy is stored (as state information) in a memory device that is accessible by the first router 18. The router then forwards the path message to the receiving access device 12 (step 208).

After it receives the path message, the receiving access device 12 determines at step 210 if it is to communicate with the sending access device 10. For example, the user on the receiving access device 12 may not have the need or desire to communicate with the sending access device 10. If the receiving access device 12 is not to communicate with the sending access device 10, then the process continues to off page connector "A." Conversely, if the receiving access device 12 is to communicate with the

sending access device 10, then the process continues to off page connector "B." Both off page connectors "A" and "B" continue on figure 2B.

From off page connector "A," the process continues to a denial process, in which a denial message is forwarded to the first router 18 (step 212), and then to the sending access device 10 (step 214). In alternative embodiments, communication is denied merely by the receiving access device 12 not responding to the path message within a preselected time interval.

From off page connector "B," the process continues to step 216, in which a RSVP ResV message ("reservation message") is forwarded to the first router 18. The first router 18 consequently applies the policy upon receipt of the reservation message (step 218). To that end, the first router 18 first matches the reservation message with its corresponding path message based upon network traffic flow. After the messages are matched, the first router 18 sends the reservation message (with the DSCP marking in a DCLASS object and/or 802.1p marking in a TCLASS object) toward the sending access device 10 (step 220). Accordingly, the policy is applied by the first network device without the further assistance of the policy server 16. Moreover, the reservation message is not processed or even received by the policy server 16.

Upon receipt of the reservation message with the applied policy, the sending access device 10 begins marking the relevant network traffic with the specified DSCP and/or 802.1p marking (step 222). Various network devices in the network 14 thus use QoS resources for the relevant network traffic flow based upon the DSCP and/or 802.1p marking.

Offloading the reservation message processing to the first router 18 reduces the processing required by the policy server 16. Accordingly, the policy server 16 may substantially simultaneously process more path messages than those policy servers that process both the path message and reservation message. Consequently, a single policy server utilizing illustrative embodiments is scalable to process more users and more network devices (e.g., routers with the described functionality of the first router 18). Of course, discussion of a router is exemplary only. For example, the first router 18 may be any network device, such as a switch or a bridge. Additionally, instead of being executed by the first router 18, the second router 20 or other router in the network 14 could execute the noted functions described in figures 2A and 2B. It is preferable, however, that a router that serves as an entry point into the network 14 for the sending access device 10 (e.g., the first router 18) be utilized to execute illustrative embodiments.

Illustrative embodiments may be implemented for a variety of network products. For example, illustrative embodiments may be utilized for massive deployment of Voice over IP, and for allowing access to the IP backbone that serves a wireless network. Protocols other than RSVP may also be used. For example, IGMP can be used as signaling for access control to a multimedia distribution network served by an IP-based backbone for digital video distribution (e.g., video on demand and digital pay per view television). Notice that when used with IGMP, the receiver of information issues the request for information. For RSVP, the sender of information issues the request to send. Illustrative embodiments permit these and other technologies to be deployed in a wide scale without being limited by the

scalability of the policy server 16 for access, billing, and other types of policies that need to be applied.

Some embodiments of the invention may be implemented in any conventional computer programming language. For
5 example, illustrative embodiments may be implemented in a procedural programming language (e.g., "C") or an object oriented programming language (e.g., "C++"). Alternative embodiments of the invention may be implemented in discrete hardware elements, and/or as preprogrammed hardware elements
10 (e.g., application specific integrated circuits and digital signal processors), or other related components.

Alternative embodiments of the invention may be implemented as a computer program product for use with a computer system. Such implementation may include a series
15 of computer instructions fixed either on a tangible medium, such as a computer readable media (e.g., a diskette, CD-ROM, ROM, or fixed disk), or transmittable to a computer system via a modem or other interface device, such as a communications adapter connected to a network over a medium.
20 The medium may be either a tangible medium (e.g., optical or analog communications lines) or a medium implemented with wireless techniques (e.g., microwave, infrared or other transmission techniques). The series of computer instructions preferably embodies all or part of the
25 functionality previously described herein with respect to the system. Those skilled in the art should appreciate that such computer instructions can be written in a number of programming languages for use with many computer architectures or operating systems. Furthermore, such
30 instructions may be stored in any memory device, such as semiconductor, magnetic, optical or other memory devices, and may be transmitted using any communications technology,

such as optical, infrared, microwave, or other transmission technologies. It is expected that such a computer program product may be distributed as a removable medium with accompanying printed or electronic documentation (e.g.,
5 shrink wrapped software), preloaded with a computer system (e.g., on system ROM or fixed disk), or distributed from a server or electronic bulletin board over the network (e.g., the Internet or World Wide Web).

Although various exemplary embodiments of the invention
10 have been disclosed, it should be apparent to those skilled in the art that various changes and modifications can be made which will achieve some of the advantages of the invention without departing from the true scope of the invention.